

A CDS Option Miscellany

R. Martin

January 4, 2012

Abstract

CDS options allow investors to express a view on spread volatility and obtain a wider range of payoffs than are possible with vanilla CDS. We give a detailed exposition of different types of single-name CDS option, including options with upfront protection payment, recovery options and recovery swaps, and also presents a new formula for the index option. The emphasis is on using the Black-76 formula where possible and ensuring consistency within asset classes. In the framework shown here the ‘armageddon event’ does not require special attention.

1 Introduction

CDS options allow investors to express a view on spread volatility and give a wider range of payoffs than are possible with vanilla CDS. Index options are commonly traded: the reader is referred to [7] for a discussion of trading strategies and hedging of index options. The market in the single-name options is not well developed, but the options are embedded in a number of structured products, for example cancellable single-name protection. With the standardisation of the single-name CDS and the move to central clearing, it is likely that single-name CDS options will become more relevant, in the same way that equity options have been around for years. Accordingly, their definition and valuation merits study.

The valuation of single-name options with running spread was originally performed in the Black-76 framework some years ago [8, 12, 13]. In this article we briefly review it and develop the subject to point out what happens when the strike is part-upfront, by which we mean that on exercise the CDS protection is to be paid wholly or partly upfront rather than all-running—

which is how contracts now trade¹. Upfronts introduce complications. The ‘Black-76 formula’ is not directly applicable because on exercise the option is not being paid for in units of risky annuity, but instead in a mixture of cash and risky annuity. We could simply start with a hazard rate model formulation, and build everything up from first principles. This is the approach taken by more recent authors such as [1, 3]. However, we would lose equivalence with Black-76 in the all-running case, even for a lognormal hazard rate model. (Lognormal hazard rate dynamics do not exactly correspond to lognormal credit spreads. Also, in the all-running case the spread is assumed for mathematical convenience to be lognormal in the survival measure with the risky PV01 as numeraire, not the money market account.) We wish to keep consistency with Black-76 as far as possible as that is the standard model used on trading desks, even if Brownian spread dynamics are ‘clearly not right’. Another complication of upfronts is that the no-knockout option² contains an embedded option on realised recovery, for which pricing models need to be developed. We deal with this too.

CDS indices are different from single-name options in that they trade no-knockout: the payoff arising from defaults in the life of the option is referred to as front-end protection (FEP). It is pointed out by Pedersen [10] that one cannot simply add the FEP to a knockout payer option to get the no-knockout price: the exercise decision depends on the accrued loss. He addresses this problem by introducing a loss-adjusted spread which he then models as lognormal. But Brigo ([4] & references therein) points out that spread-based formulae have a problem: to derive them rigorously requires changing numeraire to the associated PV01. In the ‘armageddon’ event of all names defaulting, the numeraire collapses to zero and the spread becomes undefined, so this event has to be conditioned against and treated separately. Brigo & Morini seek to estimate the probability of collapse using the tranche market, but this market has hardly been a paragon of efficiency, and they use the Gaussian copula model which did not even calibrate during the crisis: also, what is one to do with an index on which no tranches are traded (e.g. ITX.XO)? To make matters worse, they ignore in common with other authors [11] that a CDS index is not a portfolio of single names: it trades upfront with a fixed coupon and in the options the form of the strike leg payment needs careful attention (Pedersen [10] and O’Kane [9] get this

¹i.e. with fixed coupon, so that when trades are closed out, they net off with IO strip left over. This ‘Big Bang’ is necessary for the introduction of centralised clearing.

²Knockout (‘KO’) means that the option expires worthless if a credit event occurs before the option expiry; no-knockout (‘NKO’) means that it does not and can therefore be exercised into the defaulted name.

right, though O’Kane’s exposition is clearer). Here we model an index option as an option on a present value (PV), not a spread. It makes the put-call and spot-forward parity relations very clear, and the ‘armageddon problem’ disappears.

Notation. RPV01 means the PV of a risky annuity (IO strip) on the issuer; the symbol for its value on date t will be $V_t^1(T_1, T_2)$, indicating the dates T_1, T_2 between which the annuity is to be paid. The value of the annuity collapses to 0 on default. By the ‘default-PV’ of a CDS, $V_t^D(T_1, T_2)$, we mean the PV of the protection leg. As usual B denotes the riskfree discount factor (riskfree zero-coupon bond price), \mathfrak{B}_t the rolled-up money market account ($\mathfrak{B}_0 = 1$), and for single issuers τ is the default time and $Q(t)$ the survival probability. C_+, C_- generically denote a call or put on some underlying, the interpretation being obvious from the context.

2 Single-name options

2.1 Knockout, running

Until now, single-name options have been assumed to trade knockout and exercise into protection paid as a running premium. The payer price C_+ (the receiver price will be analogous and we omit it throughout to save space) is

$$\begin{aligned} C_+^{\text{ko,r}} &= \mathbf{E}_0 \left[\frac{1}{\mathfrak{B}_{t_E}} (V_{t_E}^D(t_E, T) - s_K V_{t_E}^1(t_E, T))^+ \mathbf{1}(\tau > t_E) \right] \\ &= \mathbf{E}_0 [(\bar{s}_{t_E} - s_K)^+ V_t^1(t_E, T)] \end{aligned}$$

where \bar{s} denotes the par spread, s_K the strike spread, V_t^1 the RPV01, t_E the expiry date of the option, and T the maturity date of the underlying swap (today=0). Now V_t^1 is a collapsing numeraire as it is zero after default, so we pass to the survival measure [13] \mathbf{E}^* defined as

$$\mathbf{E}^*[Z] \equiv \frac{\mathbf{E}[\mathfrak{B}_{t_E}^{-1} Z V^1(t_E, T) \mathbf{1}(\tau > t_E)]}{\mathbf{E}[\mathfrak{B}_{t_E}^{-1} V^1(t_E, T) \mathbf{1}(\tau > t_E)]},$$

note that in the denominator the survival indicator is unnecessary as $V_t^1 = 0$ if default occurs before time t . Then

$$C_+^{\text{ko,r}} = \mathbf{E}_0^* [(\bar{s}_{t_E} - s_K)^+ V_0^1(t_E, T)].$$

Under \mathbf{E}_0^* , \bar{s}_t is assumed to be lognormal, and then the expectation evaluates using Black-76 to get:

$$C_+^{\text{ko,r}} = (\bar{s}_F \Phi(d_+) - s_K \Phi(d_-)) V_0^1(t_E, T) \quad (1)$$

$$d_{\pm} = \frac{\ln(\bar{s}_F/s_K) \pm \frac{1}{2}\sigma^2 t_E}{\sigma \sqrt{t_E}}.$$

Note that \bar{s}_F (the knockout forward spread) is equal to $V_0^D(t_E, T)/V_0^1(t_E, T)$ until default, at which point it becomes undefined. When performing any statistical analysis on the credit spread, for example to estimate realised volatility, one is naturally observing it prior to default anyway, so the survival measure is the natural frame of reference.

2.2 No-knockout, running

A no-knockout option remains valid after default. By splitting it into a survival-contingent and a default-contingent part, one sees that the NKO payer must be worth the sum of the KO payer option and the FEP, so

$$C_+^{\text{nko,r}} = C_+^{\text{ko,r}} + V_0^D(0, t_E).$$

A NKO receiver option is worth no more than a knockout receiver option because a receiver option is never exercised into a defaulted credit: one would receive no running spread but pay par minus recovery. However, if the premium leg on exercise is partly paid upfront, the position is more difficult, as we are about to see.

2.3 Knockout, upfront+running

If the premium on exercise is to be paid partly (or wholly) upfront, the expression for the payer option is now

$$C_+^{\text{ko,ur}} = \mathbf{E}_0 \left[\frac{1}{\mathfrak{B}_{t_E}} (V_{t_E}^D(t_E, T) - u_K - s_K V_{t_E}^1(t_E, T))^+ \mathbf{1}(\tau > t_E) \right] \quad (2)$$

where u_K is the upfront part of the strike and s_K is the running part (in practice this is likely to be 100bp or 500bp). The RPV01 in the strike leg is as before and at maturity it depends on \bar{s}_{t_E} .

If we change numeraire in the same way as before and try to use

$$u_K + s_K V_t^1(t_E, T)$$

as numeraire, we end up with a valid option-pricing formula, but there are some problems. The main one is that it is inconsistent with the previous setup to the extent that one cannot use the same volatility for both. Of the two PV ratios

$$\left(\frac{V_0^D(t_E, T)}{s_K V_0^1(t_E, T)} \right) \text{ and } \left(\frac{V_0^D(t_E, T)}{u_K + s_K V_0^1(t_E, T)} \right),$$

(conditionally on no default) the former has higher volatility, for when the default-PV goes up, the RPV01 goes down. We therefore expect a CDS option with a strike quoted partly upfront to be worth *less* than that of an all-running one, all other things being equal. Another objection is that the second quantity is bounded, whereas the first is not, as $s_t \rightarrow \infty$, so a lognormal distribution is arguably inappropriate; however this should only be an objection for very high strikes.

To ensure consistency with (1) we therefore need to stick with lognormal spread dynamics in the V_t^1 -measure³ (which we are calling \mathbf{P}^*). However, the direct computation of the option price (2) has to be done in the \mathfrak{B}_t -measure \mathbf{P} . A change of numeraire is required. Under \mathbf{P}^* the spread is lognormal, so the expectation of some function G of the spread at time t is

$$\mathbf{E}_0^*[G(s_t)] = \int_{-\infty}^{\infty} G(\bar{s}_t(z)) \phi(z) dz, \quad \bar{s}_t(z) = \bar{s}_F e^{\sigma z \sqrt{t} - \frac{1}{2} \sigma^2 t}$$

with $\phi(z) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2} z^2)$ and \bar{s}_F denoting the forward spread. (Clearly this prescription allows the model to be embedded in a Markovian spread model, which is necessary for term structure models but not for the one-period models here.) Next, note that for any random variable Z ,

$$\mathbf{E} \left[\frac{Z}{\mathfrak{B}_t} \mathbf{1}(\tau > t) \right] \equiv \mathbf{E}^* \left[\frac{Z}{V_t^1} \right] \mathbf{E} \left[\frac{V_t^1}{\mathfrak{B}_t} \mathbf{1}(\tau > t) \right]. \quad (3)$$

We need to link the RPV01 to the spread, i.e. write V^1 as a function of \bar{s} , and to do this we use the standard RPV01 calculation Π based on a flat hazard rate curve and constant ‘assumed’ (or marking) recovery rate R . For a flat riskfree curve this is

$$\Pi(s; t, T) = \frac{1 - e^{-[r + (1+\varepsilon)s/(1-R)](T-t)}}{r + (1 + \varepsilon)s/(1 - R)} \quad (4)$$

³i.e. the survival measure with RPV01 as numeraire.

The parameter ε , which is small, will be explained presently. First, set $Z \equiv 1$ in (3) to obtain

$$\int_{-\infty}^{\infty} \frac{\phi(z) dz}{\Pi(\bar{s}_{t_E}(z); t_E, T)} = \frac{B^*(t_E)}{V_0^1(t_E, T)}$$

in which the RHS is obtained from today's CDS curve and $B^*(t_E)$ is the risky discount factor. The purpose of ε is to ensure that this relation is satisfied exactly, and a unique ε can always be found to do the job⁴. Therefore we have matched the default and premium legs of today's CDS curve, by matching the forward rate and RPV01. We now have

$$\mathbf{E}_0 \left[\frac{G(\bar{s}_{t_E})}{\mathfrak{B}_t} \mathbf{1}(\tau > t_E) \right] = \frac{\int_{-\infty}^{\infty} [G(\bar{s}_{t_E}(z); t_E, T) / \Pi(\bar{s}_{t_E}(z); t_E, T)] \phi(z) dz}{B^*(t_E) / V_0^1(t_E, T)} \quad (5)$$

and can now calculate the option price numerically from (5). This will give consistency with the Black-76 formula, because if the option exercises into all-running protection then the integral (5) gives the same result as (1).

2.4 Knockout, post-‘Big Bang’

CDS are now quoted running but trade upfront with a fixed coupon c say. This is very much like the previous section, but the RPV01 is slightly different. There is no standard convention for the options, and the payoff is best written:

$$C_+^{\text{ko,ur}} = \mathbf{E}_0 \left[\frac{1}{\mathfrak{B}_{t_E}} ((s_{t_E}^q - c) \Pi_{t_E}(s_{t_E}^q; t_E, T) - (s_K - c) \Pi_{t_E}(s_?; t_E, T))^+ \mathbf{1}(\tau > t_E) \right] \quad (6)$$

where the notation s^q indicates the quoted spread, as distinct from the par spread \bar{s} [the two are related by $(s^q - c) \Pi = (\bar{s} - c) V^1$]. The symbol $s_?$ in the strike RPV01 reflects uncertainty as to what spread to use: it could be the spot at expiry ($s_{t_E}^q$) or the strike spread (s_K). Whichever one is chosen, equation (5) can be used to value the option via a numerical integral. In this way we ensure consistency between the old-style CDS and the new-style with different coupons.

2.5 No-knockout, upfront+running

When part of the CDS premium on exercise is to be paid upfront, the no-knockout receiver option has more value than the knockout. This is because

⁴For $\varepsilon \rightarrow \infty$, $\Pi(s) \rightarrow 0$; and for $\varepsilon = -1$, $\Pi(s)$ is the riskfree PV01 which is as high as it can be.

the holder of a receiver owns a call on the realised recovery R , with strike equal to 100% minus the upfront. For example if the receiver is struck at 12% plus 500bp running, and a default occurs with 92% recovery, the option holder can exercise and make 4%. Similarly the payer has an embedded put:

$$\begin{aligned} C_+^{\text{nko,ur}} &= C_+^{\text{ko,ur}} + C_-(R; 1 - u_K) \\ C_-^{\text{nko,ur}} &= C_-^{\text{ko,ur}} + C_+(R; 1 - u_K). \end{aligned}$$

One therefore has the task of valuing the recovery option.

2.6 Recovery swaps and options

Recovery is obviously bounded to $[0, 1]$ so we wish to choose an appropriate distributional assumption. Although a Beta distribution is often suggested, we suggest the use of a Vasicek distribution instead as these are a little more tractable, in the sense that for option pricing one only needs the bivariate Normal distribution as opposed to the incomplete Beta function. Also the use of the Vasicek distribution ties in neatly with probit modelling of recovery rates. The recovery R is modelled as a transformation of a Normal variable, thus:

$$R(Z) = \Phi\left(\frac{a + bZ}{\sqrt{1 - b^2}}\right), \quad Z \sim N(0, 1)$$

which has mean $\Phi(a)$ and variance $\Phi_2(a, a; b^2) - \Phi(a)^2$. (Here $\Phi_2(x, y; \rho)$ denotes as usual the cumulative bivariate Normal distribution.) For small b the standard deviation is roughly $\phi(a)b$, which follows from the tetrachoric series expansion of Φ_2 . This enables the distribution to be parametrised in terms of understandable numbers. Notice also that, one can easily model correlated recovery rates of several issuers, if desired, just by correlating their Z -variables through a multivariate Normal. Figure 1 shows the density for mean 20% ($a = -0.842$) and three different values of the width parameter b (0.5, 0.6, 0.7).

The payoffs of a recovery swap (or lock), call, and put are, respectively, $R - K$, $(R - K)^+$ and $(K - R)^+$, at expiry, where K is the recovery strike. There is no payment if no default occurs. Using the result

$$\mathbf{E}\left[\Phi\left(\frac{a + bZ}{\sqrt{1 - b^2}}\right) \mathbf{1}(Z > \xi)\right] = \Phi_2(-\xi, a; b)$$

we can obtain the call and put payoffs with strike u as

$$\begin{aligned} \mathbf{E}[(R - u)^+] &= \Phi_2(a, c; b) - \Phi(c)u \\ \mathbf{E}[(u - R)^+] &= \Phi(-c)u - \Phi_2(a, -c; -b) \end{aligned}$$

with $c = (a - \Phi^{-1}(u)\sqrt{1 - b^2})/b$. That put-call parity is satisfied is evident from elementary symmetries in Φ_2 .

The option values $C_{\pm}(R; u)$ are these expected payoffs multiplied by the difference between the risky and riskfree discount factors (i.e. $B(t_E) - B^*(t_E)$), because the cashflow occurs at time t_E and is contingent upon default before option expiry. The recovery swap PV is obviously the difference between the strike and market recovery level, multiplied by the same factor.

In practice one obviously has to estimate the two parameters a and b . Recovery locks do trade so there is some guide from the market as to where recovery is likely to be, but there is no information about uncertainty in recovery. As an example, take the defaults in CDX.HY9 (fourteen at the time of writing). The average recovery has been 17.5% but the dispersion is very wide: the highest was Tembec with 83%, there were a couple at supposedly investment-grade levels (Quebecor 41.25%, Lear 38.4%), but the lowest were almost zero (Tribune 1.5%, Idearc 1.75%, Charter 2.375%, Visteon 3%, Abitibi 3.25%, RHD 4.875%). Thus without any knowledge of the credit one should probably use a high dispersion parameter ($b \sim 0.75$ to fit these fourteen points⁵). Of course, there is generally less uncertainty in the market over a particular traded credit, particularly at short horizon, and then the b parameter is selected so as to represent the analysts' uncertainty.

One might ask whether it is necessary to 'correlate recovery with default'. First one has to ask what this means. Importantly, the random variable Z is only observed when default occurs, so one only wants its distribution conditional on default: in any other state of the world its value is irrelevant. Its distribution may be time-dependent (so that one has to use a different a, b for different option maturities)—but use of different parameters for different maturities, most notably of course the Black-Scholes volatility, is standard practice.

For multivariate models, the question of correlation takes on a different form, and this is not relevant to the pricing of single-name options. One obviously wishes to correlate default events and there is evidence that when default rates *en bloc* are higher, recoveries are lower [5, 6]. Incidentally the incorporation of this effect is also valuable in modelling of CDO tranches as it is a convenient device for pushing losses up into the senior tranches, a perennial difficulty for tranche modellers. Now, suppose that at a particular horizon one wishes to correlate defaults and recoveries for many credits. We want to ensure that we can consistently model correlated defaults and recoveries, having specified that the distribution of recovery (conditional, as

⁵By Kolmogorov-Smirnov.

said before, on default) is Vasicek. It turns out that this is possible without restricting the dependence structure, and this is discussed in the Appendix.

2.7 Example

We consider option pricing with the following parameters: vol 100%, recovery 20%, pricing date 09-Nov-09, option expiry 20-Mar-10, swap maturity 20-Dec-14. For no-knockout options we need the recovery rate volatility, and this is given via the ‘ b ’ parameter which we take to be 0.6.

Figure 2 is for spot=500bp (forward-starting RPV01 3.723). The left-most points on the graph are the payer and receiver premiums for all-running strike (500bp), which can be checked against the standard Black-Scholes price. The other points show the prices of options in which the strike is quoted partly upfront: the conversion is that 100bp running corresponds to 3.72% upfront, and the all-upfront strike is 18.6%. As expected the option prices decrease as the proportion of the upfront increases. The KO payer should be 35bp cheaper for all-upfront strike than for all-running, and the receiver about 100bp cheaper. The NKO payer options are obviously more expensive than their KO counterparts, by an amount that decreases as the upfront part is raised. For the receiver the embedded recovery call option is too far out of the money to have noticeable value.

Figure 3 is for spot=2000bp (forward-starting RPV01 2.326). The left-most points on the graph are the payer and receiver premiums for all-running strike (this time 2000bp). Again the other points are for part-upfront strike, and the all-upfront equivalent strike is 46.5%. The difference in price between the all-upfront and all-running options is now considerably higher.

3 Index options

Index options are best treated as a separate asset class from single-name options because a literal model of an index option would require a model of correlated spreads and defaults between all the issuers. By convention they trade no-knockout, so that exercise is into the ‘original’ index without removing any credits that have defaulted in the life of the option, and this is the case that we analyse here. As pointed out by O’Kane [9]⁶, the decision to exercise or not depends on the accrued loss in the pool, as well as the

⁶This is the most coherent treatment of CDS options we have seen and the reader is strongly advised to work through his numerical examples.

current spread and RPV01, so it is not possible to simply add the front-end protection to the payer price in the way that is done for single-name options.

On exercise of an index payer, the following payments are made:

- Option buyer receives par minus recovery on all defaulted names since the option was struck;
- Option buyer effectively exercises into an index contract, at the prevailing market spread, on the remaining notional, i.e. the original notional multiplied by the proportion of names that are still undefaulted (this is usually cash-settled though);
- Option buyer pays a cash amount equal to the strike spread minus the index coupon, multiplied by a PV01 defined by convention as $\Pi_{t_E}(s_K; t_E, T)$, that is, the RPV01 assuming a flat curve at the *strike* spread (not the prevailing market spread, which would be undefined in the armageddon event), on the *full* notional.

Symbolically, we have that the value of an index payer is

$$C^{+\text{idx}} = \mathbf{E}_t \left[\frac{\mathfrak{B}_t}{\mathfrak{B}_{t_E}} \left(L_{t_E} + \frac{N_{t_E}}{N_{00}}(s_{t_E} - c)\Pi_{t_E}(s_{t_E}; t_E, T) - \frac{N_0}{N_{00}}(s_K - c)\Pi_{t_E}(s_K; t_E, T) \right)^+ \right]$$

Here $\Pi_t(s; T_1, T_2)$ is the ‘flat-CDS-curve’ PV01 calculation alluded to above⁷. It is simply a formal calculation to get from a spread to an upfront for use in settlement, and does not, for example, depend on the number of defaults in the index (though it does depend on an assumed recovery rate which is fixed by market convention), and nor is it a tradable asset as such. Also L_t is the accrued loss through defaults⁸ from time 0 to time t and N_t is the number of survived names at time t . The option is assumed to have been struck at time 0, and N_{00} denotes the number of names in the original index: for example we might have $N_{00} = 125$, $N_0 = 124$ and $N_t = 122$, meaning that the index originally had 125 names, one default occurred, then the option was written, then another two occurred, and now it is time to value the option⁹. It is also to be understood that s_t is the spread of the current version of the index.

A minor point worth ironing out is what happens when the strike spread is made arbitrarily high. According to O’Kane, the payer option should

⁷Except that eq. (4) assumes a flat hazard rate, which is not quite the same thing; and one should use the proper riskfree curve too. In essence, this means reproducing what the Bloomberg CDSW page implicitly uses.

⁸Assuming interest is paid from the auction payout, but this is a minor matter.

⁹ N_t/N_{00} is called the ‘factor’.

become worthless in that limit, on the grounds that one is paying an ‘infinite spread’ on exercise, whereas the simple procedure of taking a knockout payer Black-76 style and adding on the front-end protection causes the payer value to tend to the value of the front-end protection. Actually, both are wrong, because the exercise *premium* does not tend to infinity: its limit is, as can be easily verified,

$$\lim_{s_K \rightarrow \infty} (s_K - c)\Pi_t(s_K; t_E, T) = 1 - R$$

where R is the index recovery used in the RPV01 calculation.

We want to explicitly model the option as one on the difference between two PV’s X and Y say, i.e. the default-PV of the underlying swap and the PV of the coupon leg that would be exercised into at the prescribed strike, change numeraire to Y and then use Black-76. Note that in the context of the single names we avoided this approach for the part-upfront options, because of inconsistency with the all-running case: with the index options there is only one product to model, so we can deal with that as the *de facto* standard.

However, we can only use this construction if X and Y are positive a.s. Therefore we have to be a little careful in grouping the terms, and the following construction achieves this in an intuitive way:

$$X_t = L_t + \frac{N_t}{N_{00}}(s_t - c)\Pi_t(s_t; t, T) + \frac{cN_0}{N_{00}}A_t \quad (7)$$

$$Y_t = \frac{N_0}{N_{00}}\left((s_K - c)\Pi_t(s_K; t_E, T) + cA_t\right) \quad (8)$$

$$A_t = -\int_{t'=0}^t \frac{N_{t'}}{N_0}d\left(\frac{\mathfrak{B}_t}{\mathfrak{B}_{t'}}\right) + \int_{t'=0}^t d\left(\frac{\mathfrak{B}_t}{\mathfrak{B}_{t'}}\right) + \Pi_t(0; t, T) \quad (9)$$

The extra terms that have been added are explained as follows. The expression X_t is the PV at time t of a long index protection trade entered into at time 0 for an investor who is assumed to own a riskfree annuity to pay the cost of carry during the life of the trade, thereby making the trade fully-funded. This is because the first term in (7) is the accrued payout from defaults, the second is the PV of the remaining index protection, and in A_t the first term is the accrued carry (which steps down as defaults occur), the second is the accrued carry from the riskfree annuity, and the third is its remaining value. It is easy to see that $X_t > 0$ a.s., because the first term is nonnegative, the second term in A_t exceeds the first, and the last term in A_t exceeds the negative part of the second term in X_t . In fact X_t approaches 0 in the situation where $N_t = N_0$ (so $L_t = 0$) and $s_t \rightarrow 0$. It is also clear

that $Y_t > \frac{N_0}{N_{00}} s_K \Pi_t(s_K; t_E, T)$ a.s. by similar reasoning. Notice that all the correction terms result from the presence of the index coupon. Finally, X_t and Y_t are well-defined even in the armageddon event, for although s_t and hence $\Pi_t(s_t; t, T)$ are no longer defined, the term that references them is being multiplied by N_t/N_{00} which is zero and so it vanishes.

We need the discounted expectations of X_{t_E} and Y_{t_E} at time t , i.e.

$$\tilde{X}_t := \mathbf{E}_t[X_{t_E} \mathfrak{B}_t / \mathfrak{B}_{t_E}], \quad \tilde{Y}_t := \mathbf{E}_t[Y_{t_E} \mathfrak{B}_t / \mathfrak{B}_{t_E}].$$

By construction X_t is the PV of a sum of tradable assets, so

$$\tilde{X}_t = X_t.$$

At initiation one has simply

$$\tilde{X}_0 = \frac{N_0}{N_{00}} \left(s_0 \Pi_0(s_0; 0, T) + c(\Pi_0(0; 0, T) - \Pi_0(s_0; 0, T)) \right).$$

Not all parts of Y_t are tradable assets¹⁰. However, the expectation of Y_t can be approximated easily enough:

$$\tilde{Y}_t \approx \frac{N_0}{N_{00}} \left((s_K - c) \Pi_t(s_K; t_E, T) + c \mathbf{E}_t \left[\int_{t'=0}^{t_E} \left(1 - \frac{N_{t'}}{N_0} \right) d \left(\frac{\mathfrak{B}_t}{\mathfrak{B}_{t'}} \right) \right] + c \Pi_t(0; t_E, T) \right). \quad (10)$$

The expectation in the second term, which is the writedown of the risky index annuity caused by defaults, is approximately $\frac{1}{2}(1 - Q_t(t_E))$, where $Q_t(\cdot)$ is the index survival probability (expected proportion of defaulted names, as seen from time t), so at time 0 we have

$$\tilde{Y}_0 \approx \frac{N_0}{N_{00}} \left(s_K \Pi_0(s_K; t_E, T) + c \left(\frac{1}{2}(1 - Q_0(t_E)) + \Pi_0(0; t_E, T) - \Pi_0(s_K; t_E, T) \right) \right).$$

The final result is then:

$$C_t^{+\text{idx}} = \tilde{X}_t \Phi(d_+) - \tilde{Y}_t \Phi(d_-), \quad (11)$$

$$\tilde{X}_t = \text{Eq. (7)}; \quad \tilde{Y}_t = \text{Eq. (8)}; \quad d_{\pm} = \frac{\ln(\tilde{X}_t / \tilde{Y}_t) \pm \frac{1}{2} \sigma^2 (t_E - t)}{\sigma \sqrt{t_E - t}}. \quad (12)$$

¹⁰The first is not strictly a tradable annuity and technically we should worry about a convexity correction arising from its interest-rate risk, but this is of too little importance in context; the third and fourth combined obviously are tradable, and the second is the accrued index carry which needs to be estimated.

The delta is given, as usual, by $\Phi(d_+)$. We have deliberately left the relevant quantities as PV's rather than converting them back to spreads. One could try to write X_t and Y_t in terms of a corrected forward spread and a corrected strike (the corrections being for the coupon effect and the accrued loss), viz:

$$\begin{aligned}\tilde{X}_t &= s_{F,t}^* \Pi_t(s_t; t, T), & s_{F,t}^* &= \frac{N_t}{N_{00}} s_0 + \frac{L_t}{\Pi_t(s_t; t, T)} + c(\dots) \\ \tilde{Y}_t &= s_K^* \Pi_t(s_K; t_E, T), & s_K^* &= s_K + c(\dots).\end{aligned}$$

But this does not seem to achieve anything useful: the PV01s are different, and it also unearths the previously-buried problem of what happens in the armageddon event (\tilde{X}_t is well-defined but $s_{F,t}^*$ is not). Notice incidentally that some of the adjusted forward corrections in the literature¹¹ can result in a negative adjusted forward and/or strike spread, in the admittedly unlikely situation of the strike being very low. The construction used here renders this impossible.

Incidentally the CDX.HY (high-yield) index is quoted as a bond price, but still trades as a swap. (For example, a price of 97.625 means that a buyer of index protection pays 2.375% upfront plus 500bp running.) It is convenient to work with the prevailing upfront U_t , and an upfront strike u_K ; these will be negative if the associated bond price exceeds 100. The index payer, which is usually referred to as a put, has value

$$\mathbf{E}_t \left[\frac{\mathfrak{B}_{t_E}}{\mathfrak{B}_t} \left(L_t + \frac{N_t}{N_{00}} U_t - \frac{N_0}{N_{00}} u_K \right)^+ \right].$$

Our construction is now the same as before, only simpler:

$$\begin{aligned}X_t &= L_t + \frac{N_t}{N_{00}} U_t + \frac{cN_0}{N_{00}} A_t \\ Y_t &= \frac{N_0}{N_{00}} (u_K + cA_t)\end{aligned} \tag{13}$$

and A_t as before (9). In fact, all one is doing here is pricing a bond option on the assumption that the asset-swap spread of the underlying bond, which is linear in the bond price, is lognormal¹². Implicitly we thereby permit losses

¹¹e.g. O'Kane [9, §11.7, Eq. 11.9]

¹²As the price distribution is (almost) always negatively skewed, and the ASW spread positively, a lognormal ASW spread seems a better idea than a lognormal bond price. One could model the yield spread (I- or Z- spread) as lognormal, which in principle is the best solution as those are unbounded, but it introduces convexity corrections.

that exceed the notional of the index, with tiny probability. We do not regard this as important, and it simply alters the shape of the volatility smile that invariably accompanies lognormal models. Analogously in the Vasicek or Hull-White interest rate models with Gaussian spot rates, negative rates are permitted, but they are so unlikely as to be of no practical concern.

The treatment of index options that we have given here is not revolutionary, but rather is intended to deal with the payout more carefully and intuitively. As we have been at pains to work in PVs rather than spreads, the reader might justifiably be worried that our volatility parameter σ is incompatible with the quoted (spread) volatility. For investment-grade indices this is not a problem and they can be used interchangeably, but for high-yield they cannot: to reprice a contract one needs a lower ‘PV vol’ in the equations here than the spread vol seen on dealer runs. This is because the par CDS spread, seen on the runs, is a convex function of the PV. As options are always quoted in price terms (bp upfront)—mainly as a result of no standard formula—with the vol being used merely as a guide, this is unlikely to result in difficulties as long as it is appreciated.

3.1 Example

Figure 4(a) shows the results for CDX.IG13 from a recent dealer run (initiation and pricing date 09-Nov-09; option expiry 20-Mar-09; index maturity 20-Dec-14; coupon 100bp; spot 100bp). The vols are inferred from the option prices and compared with the quoted vols. The discrepancy is pretty small (2–3 vol points; equivalently about 2bp on the ATM price). Figure 4(b) shows the results for CDX.HY13 on the same date (coupon 500bp; spot price 93.25). As a complication, one credit event had occurred but its auction had not, so the index was trading with a 100-name default leg but a 99-name coupon leg (thus $\frac{N_0}{N_{00}} = 0.99$ but $L_0 = 0$). The vols are now considerably different, and it is intriguing to note that ours have less ‘smile’.

4 Conclusions

We have shown how to value, in the Black-76 framework, single-name CDS options in which the strike is quoted wholly or partially upfront and shown that the proportion of upfront has a significant impact on price, if consistency with the standard ‘all-running’ case is to be achieved. We have also pointed out that no-knockout single-name options that are quoted wholly or partially upfront contain an embedded recovery option, and given a simple explicit formula for the price. For the index options we have sharpened the

existing literature by taking more care over the PV of the strike leg; the resulting formula is in essence the Black-76 price of a fully-funded index contract that is assumed to have a lognormal PV.

5 Appendix

We discuss the correlation of recovery and default in a multivariate setting (see e.g. [14]). This goes well beyond the scope of this paper and is only intended as a brief justification of our previous assertion that the univariate model we have given can be extended to the multivariate case (in infinitely many ways, in fact).

This is most conveniently done with the assistance of a risk-factor, i.e. a random variable \mathcal{A} say (which need not be univariate). Conditionally on \mathcal{A} , all recoveries and defaults are completely independent. Let a particular credit have a conditional default probability $p(\mathcal{A})$, so that $\mathbf{E}[p(\mathcal{A})] = \bar{p}$ (the average default probability, in this case obtained from the CDS market). Let F be an arbitrary (cumulative) distribution function, and β be a coefficient that will couple the recovery variable Z to the risk-factor \mathcal{A} . Define

$$F_{\sharp}(z) := \frac{1}{\bar{p}} \mathbf{E}[F(z - \beta \cdot \mathcal{A})p(\mathcal{A})].$$

Then define the conditional distribution of Z on \mathcal{A} to be

$$\mathbf{P}[Z < z | \mathcal{A}] = F(F_{\sharp}^{-1}(\Phi(z)) - \beta \cdot \mathcal{A}).$$

The distribution of Z conditionally on default of the issuer within the horizon t_E , but unconditionally on \mathcal{A} , is

$$\begin{aligned} \mathbf{P}[Z < z | \tau < t_E] &= \frac{1}{\bar{p}} \mathbf{E}[\mathbf{P}[Z < z \cap \tau < t_E | \mathcal{A}]] \\ &= \frac{1}{\bar{p}} \mathbf{E}\left[p(\mathcal{A})F(F_{\sharp}^{-1}(\Phi(z)) - \beta \cdot \mathcal{A})\right] = F_{\sharp}(F_{\sharp}^{-1}(\Phi(z))) = \Phi(z) \end{aligned}$$

which is what we wanted it to be. This procedure of constructing what might be described as ‘additive factor-copulas’ is quite generally applicable and shows that one can always match a given marginal distribution. As stated above, β plays the role of a correlation parameter, while the distribution chosen for F is completely arbitrary¹³.

The author acknowledges numerous valuable discussions with Philipp Schönbucher. Email: rmartin@ahl.com

¹³As long as it is continuous. The Gaussian gives, of course, the Gaussian copula.

References

- [1] H. Ben-Ameur, D. Brigo, and E. Errais. A dynamic programming approach for pricing CDS and CDS options. www.defaultrisk.com, 2006.
- [2] T. Bounhar and L. Luciani. Two generic frameworks for credit index volatility products and their application to credit index options. www.defaultrisk.com, 2008.
- [3] D. Brigo and N. El-Bachir. An exact formula for default swaptions pricing in the the SSRJD stochastic intensity model. www.defaultrisk.com, 2008.
- [4] D. Brigo and M. Morini. Last option before the armageddon. *RISK*, 22(9):118–123, 2009.
- [5] M. Bruche and C. González-Aguado. Recovery rates, default probabilities, and the credit cycle. www.defaultrisk.com, 2008.
- [6] P. Dobranszky. Joint modelling of CDS and LCDS spreads with correlated default and prepayment intensities and with stochastic recovery rate. www.defaultrisk.com, 2008.
- [7] S. Doctor and J. Goulden. An introduction to credit options and credit index volatility. *JP Morgan Credit Derivatives Research*, March, 2007.
- [8] J. Hull and A. White. Valuing credit default swap options. *J. Derivatives*, 10(3):40–50, 2003.
- [9] D. O’Kane. *Valuing single-name and multi-name credit derivatives*. Wiley, 2008.
- [10] C. Pedersen. Valuation of portfolio credit default swaptions. *Lehman Brothers Quantitative Credit Research*, November, 2003.
- [11] M. Rutkowski and A. Armstrong. Valuation of credit default swaptions and credit default index swaptions. www.defaultrisk.com, 2008.
- [12] P. Schönbucher. *Credit Derivatives Pricing Models*. Wiley, 2003.
- [13] P. Schönbucher. A measure of survival. *RISK*, 17(8):79–85, 2004.
- [14] D. Tasche. The single risk factor approach to capital charges in case of correlated loss given default rates. www.defaultrisk.com, 2004.

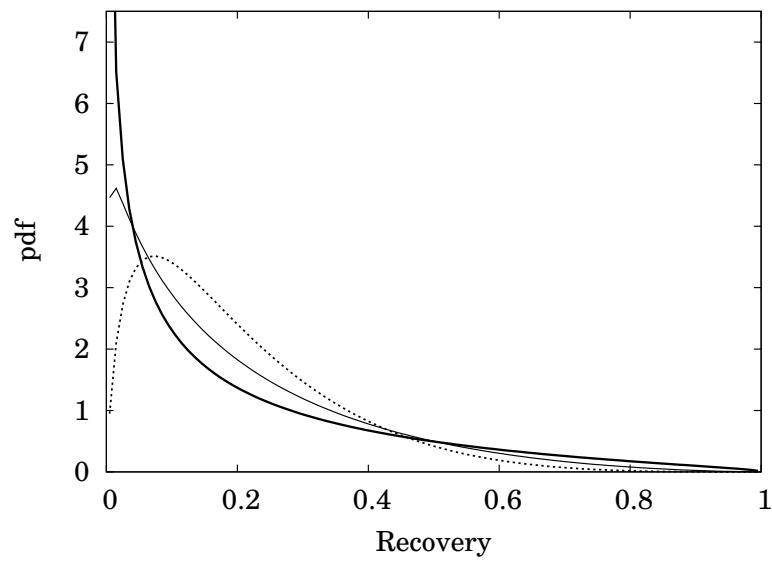


Figure 1: Vasicek distributions with mean 0.2 and width parameter 0.5 (dotted line), 0.6 (solid, light), 0.7 (solid, heavy).

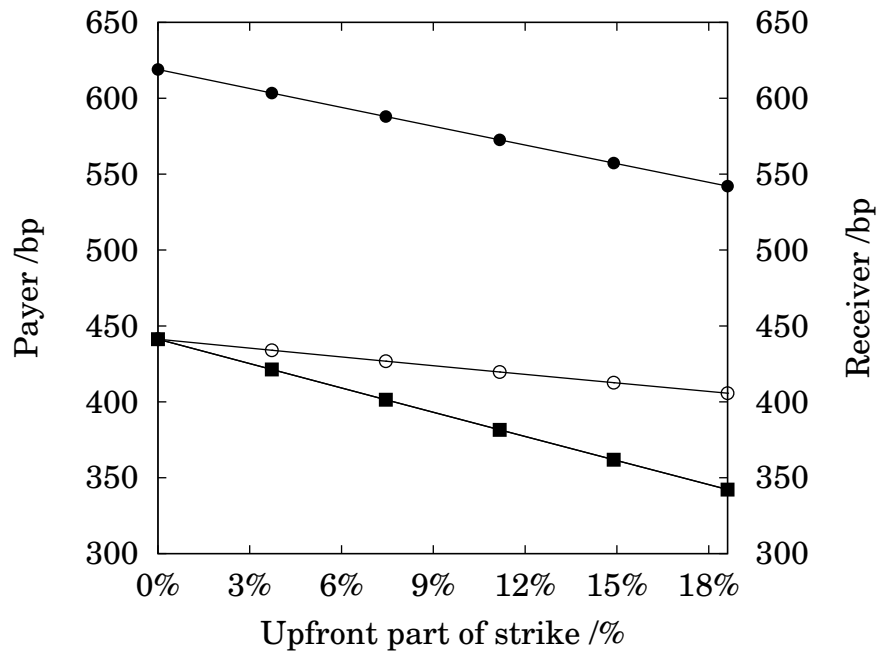


Figure 2: Knockout and no-knockout option prices as a function of the upfront part of the strike; spot 500bp. Payers: \circ =KO, \bullet =NKO, both on LH axis. Receivers: \square =KO, \blacksquare =NKO, both on RH axis (they differ by $< 0.1\text{bp}$). Running part of strike decreases linearly from 500bp to 0bp.

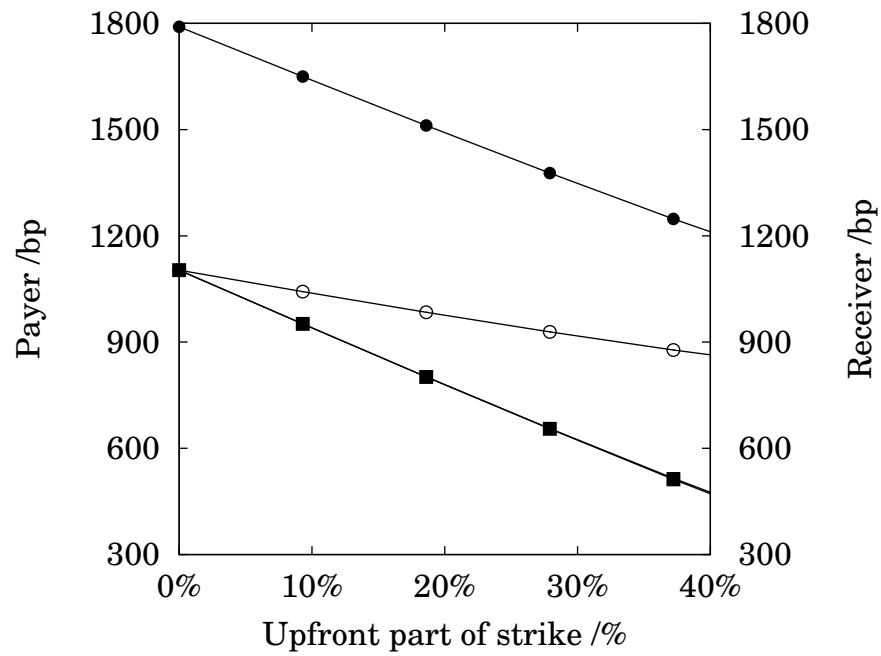
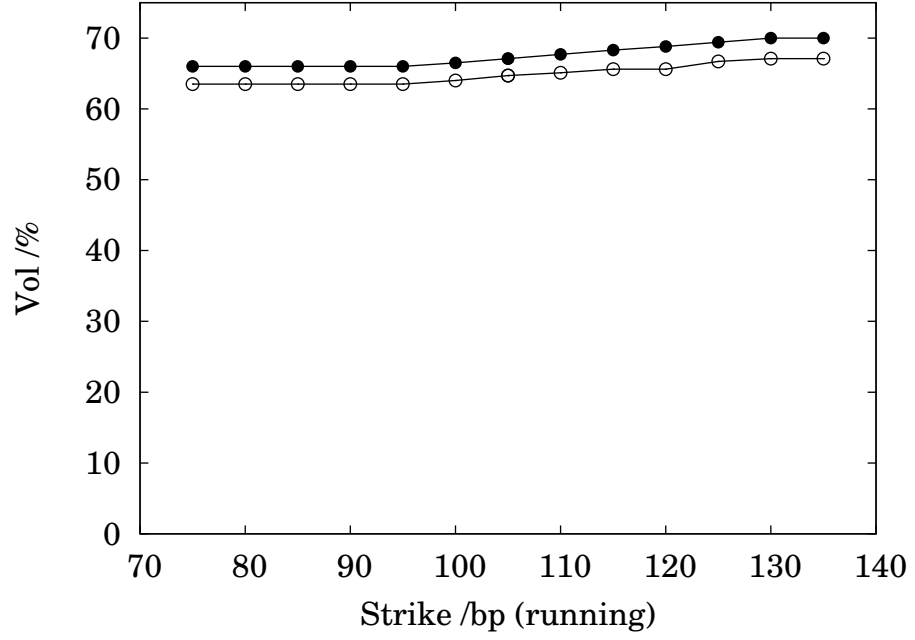
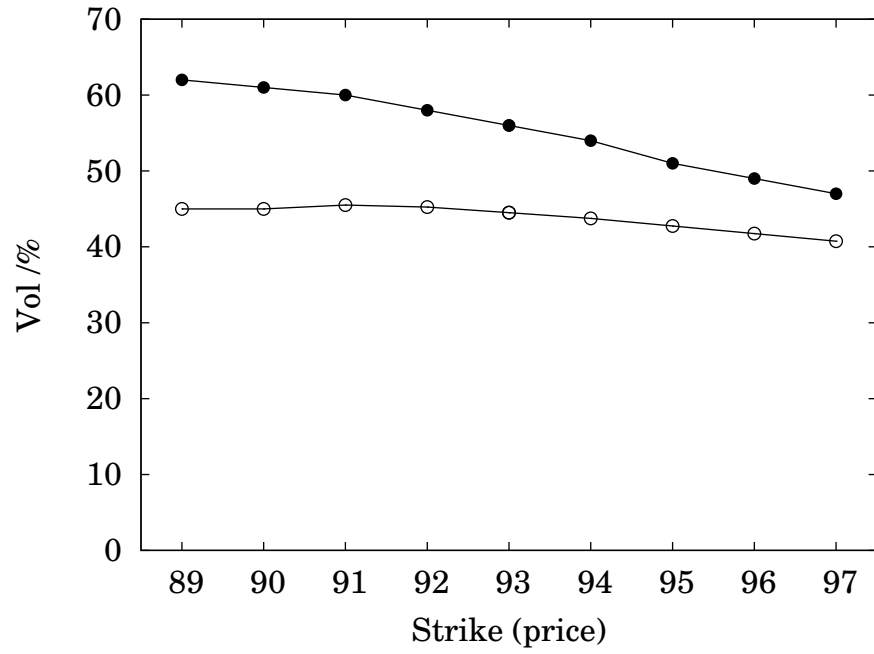


Figure 3: Knockout and no-knockout option prices as a function of the upfront part of the strike; spot 2000bp. Payers: \circ =KO, \bullet =NKO, both on LH axis. Receivers: \square =KO, \blacksquare =NKO, both on RH axis. Running part of strike decreases linearly from 2000bp to 0bp.



(a)



(b)

Figure 4: Implied vols for (a) CDX.IG13, (b) CDX.HY13. ●=quoted, ○=ours (to reproduce option price)